

Genetic Modeling of Battery Performance

10 July 2009

Albert H. Zimmerman
Electronics and Photonics Laboratory
Physical Sciences Laboratories

Prepared for:

Space and Missile Systems Center
Air Force Space Command
483 N. Aviation Blvd.
El Segundo, CA 90245-2808


Authorized by: Engineering and Technology Group

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED

This report was submitted by The Aerospace Corporation, El Segundo, CA 90245-4691, under Contract No. FA8802-09-C-0001 with the Space and Missile Systems Center, 483 N. Aviation Blvd., El Segundo, CA 90245. It was reviewed and approved for The Aerospace Corporation by B. Jaduszliwer, Principal Director, Electronics and Photonics Laboratory; and D. C. Marvin, Principal Director, Research and Program Development Office. David E. Davis was the project officer for the Mission-Oriented Investigation and Experimentation (MOIE) program.

This report has been reviewed by the Public Affairs Office (PAS) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication. Publication of this report does not constitute Air Force approval of the report's findings or conclusions. It is published only for the exchange and stimulation of ideas.


David E. Davis
SMC/EA

REPORT DOCUMENTATION PAGEForm Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)

10-07-2009

2. REPORT TYPE**3. DATES COVERED (From - To)****4. TITLE AND SUBTITLE**

Genetic Modeling of Battery Performance

5a. CONTRACT NUMBER

FA8802-09-C-0001

5b. GRANT NUMBER**5c. PROGRAM ELEMENT NUMBER****6. AUTHOR(S)**

Albert H. Zimmerman

5d. PROJECT NUMBER**5e. TASK NUMBER****5f. WORK UNIT NUMBER****7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**The Aerospace Corporation
Physical Sciences Laboratories
El Segundo, CA 90245-4691**8. PERFORMING ORGANIZATION
REPORT NUMBER**

TR-2009(8550)-9

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)Space and Missile Systems Center
Air Force Space Command
483 N. Aviation Blvd.
El Segundo, CA 90245**10. SPONSOR/MONITOR'S ACRONYM(S)**
SMC**11. SPONSOR/MONITOR'S REPORT
NUMBER(S)****12. DISTRIBUTION/AVAILABILITY STATEMENT**

Approved for public release; distribution unlimited.

13. SUPPLEMENTARY NOTES

20100310194

14. ABSTRACT

A state-of-the-art, first-principles lithium-ion cell model has been used to describe the changes in the performance of a 50-Ah cell during the first 7563 cycles of its life. This model was fit to the life test database using a genetic adaptation algorithm that provided degradation trends for nine different parameters that described the processes most likely to undergo change during the life of the cell. Based on the model trending, most of the capacity degradation seen in this cell arises from the net loss of lithium that can reversibly cycle between the anode and cathode. The model also indicates that increasing lithium diffusion resistance and increasing surface layer polarization on the cathode are contributing to the diminished chargeability of the cell, along with the observed decrease in cathode capacity. Thus, the cathode appears to be changing more than the anode during the cycling.

15. SUBJECT TERMS

Battery, Lithium Ion, Model, Genetic Analysis, Capacity Loss, Life Test, Adaptive Modeling, Life Prediction

16. SECURITY CLASSIFICATION OF:**a. REPORT**

UNCLASSIFIED

b. ABSTRACT

UNCLASSIFIED

c. THIS PAGE

UNCLASSIFIED

**17. LIMITATION
OF ABSTRACT**

Leave blank

**18. NUMBER
OF PAGES**

29

**19a. NAME OF RESPONSIBLE
PERSON**

Albert Zimmerman

**19b. TELEPHONE NUMBER
(include area code)**

(310)336-7415

Contents

1.	Introduction	1
2.	Model Description.....	3
3.	Genetic Adaptation to Test Data	9
3.1	Genetic Algorithm for Modeling Li-Ion Cells	10
3.2	Test Data Used for Modeling	12
3.3	Adaptive Modeling Approach	14
4.	Analysis of Model Results.....	17
4.1	Initial Cell State of Charge	17
4.2	Cathode Capacity	18
4.3	Lithium-Ion Diffusion Rate.....	18
4.4	Cathode Surface Polarization Time Constant	20
4.5	Cathode Surface Polarization Voltage.....	20
4.6	SEI/Electrolyte Resistance	21
4.7	Anode Capacity	21
4.8	Anode Precharge	23
4.9	Cathode Charge Transfer Resistance.....	23
4.10	Life Projection Based on Trend Analysis.....	24
5.	Conclusions	29

Figures

1.	Finite elements used to model the macroscopic spatial dimensions of a prismatic lithium-ion battery cell.	3
----	--	---

2. Electrical circuit showing electronic and ionic conduction, as well as charge transfer through the cell model.....	4
3. Illustration of the charge transfer processes modeled in the cathode active material.	5
4. Illustration of the polarization layer required at the cathode active material to enable ionic transfer to the electrolyte in Li-ion cells.....	6
5. Illustration of the processes involved in the anode energy storage materials.....	7
6. Typical voltage and current behavior during a 40% depth-of-discharge cycle.	12
7. Typical cell and base plate thermal profile during the 40% depth-of-discharge cycling...	13
8. Capacity cycles used to model Li-ion cell changes during life testing.....	14
9. Distribution of initially selected and optimized populations for two of the nine parameters adjusted to model cell performance changes over life.	15
10. Changes in average deviation between the model and the data during the 3360 generations involved in the genetic optimization illustrated in Figure 9.	16
11. Measured and optimized model cell voltages after 5505 cycles.	16
12. Model trend in initial cell state of charge during cycling.	17
13. Model trend in total cathode capacity for lithium during cycling.	18
14. Model trend in the lithium-ion diffusion resistance during cycling.	19
15. Model trend in the cathode surface polarization time constant during cycling.	20
16. Model trend in the cathode surface polarization voltage during cycling.	21
17. Model trend in the SEI resistance (ohm cm^2) during cycling.	22
18. Model trend in the total anode capacity for lithium during cycling.	22
19. Model trend in the anode precharge during cycling.	23
20. Model trend in the cathode charge transfer resistance during cycling.	24
21. Model projection of cycles until the end-of-discharge voltage falls to 3.0 V after 42,050 cycles	25
22. Plot of the remaining cell capacity as a function of the square root of the number of cycles completed at 40% depth of discharge	26
23. Capacity trend when plotted as a linear function of cycles.	26

Table

1. Initial Adaptation Ranges and Optimized Values for Parameters Adjusted in the Cell Model to Describe Changes in Cell Performance During Life Test	14
---	----

1. Introduction

The long-term performance changes that occur during years of cycling rechargeable batteries are often caused by a large number of gradual changes in the internal chemical state of the battery cells. Understanding the origin of these performance changes is critical to improve battery cell designs, as well as to predict the reliability and life of the batteries in critical applications where the batteries cannot be readily replaced or serviced, such as in satellites. For lithium-ion battery cells, the performance changes are often lumped into two easily measurable categories: capacity loss and resistance increase. These two categories capture the degradation empirically; however, they do not enable the physical processes that control the performance changes to be readily determined.

Lengthy life testing of batteries is commonly used to verify their life capability for satellite applications. Such life testing, while accurately evaluating the performance changes that occur in the battery, similarly does not generally provide a clear understanding of the physical and chemical changes within the battery cells that are responsible for the degradation in performance. Destructive physical analysis of cells can provide such information on degradation or failure modes at the end of the life test. However, it is frequently desirable to understand what types of degradation are occurring long before reaching the end of a life test that could take 10 years or longer to run to completion, or to understand how battery life would be affected by test conditions that differed from those applied to the cells during the life test. These kinds of questions generally cannot be answered solely through the combination of life testing and destructive physical analysis.

A first principles model of a lithium-ion battery cell, if it accurately includes all the key processes that influence performance, is capable of correlating observed performance changes during life tests with changes in the chemical or physical state of the electrodes or other components internal to the battery cell. However, a first principles model that accurately describes the performance of a lithium-ion cell during its lifetime must include numerous processes, which overall are defined by dozens of parameters that can potentially change over the lifetime of the cell. Thus, the identification of the model parameter changes that are most consistent with actual life test performance signatures is the best method for applying and using a first principles battery cell model.

Unfortunately, it is typically difficult to fit a high-fidelity first principles battery cell model to a life test database. The model typically contains 10–20 parameters that can potentially change over a relatively wide range as the cells age. The effects of these parameters are frequently non-linear and interact in terms of affecting cell performance, thus requiring highly iterative regression analyses to fit the model to the life test data. The computer time that is required for such a regression approach can be prohibitively long. Another significant obstacle can arise from the local minima that are frequently present in sophisticated battery cell models, and which can trap a regression analysis into a solution that is not optimum. Local minima frequently arise in models that use finite difference approximations in either the spatial dimensions that describe the cell, or use time steps to describe the temporal development of the cell. Local minima effects may be less severe in models that are based on contin-

uum temporal or spatial mechanics, but computational complexity can often limit the physical and chemical fidelity of such continuum models.

Here, a technique is described that enables the chemical and physical changes in battery cells to be determined from their life test data. The technique uses a first principles model that, in general, accurately describes how battery cell performance is affected by the possible chemical processes that can occur over time within the cell. A genetic adaptation algorithm is used to adapt the model to the observed life test data as the batteries are cycled. At any point in time during the life test, the adapted model shows the specific changes in the cells that are responsible for their degradation based on the best fit to the chemical and physical processes included in the model.

The finite-element modeling methodology used will be described here, followed by a description of the optimized genetic algorithm used to adapt the model to a specific life test database. The capability of the genetic algorithm to handle local minima while exploring wide regions in the fitting surface will be discussed, along with a comparison of its computational time relative to the speed of traditional iterative regression methods. An example will finally be provided and discussed that illustrates how the technique can be used with a lithium-ion cell life test database to determine the causes for the observed degradation in cell performance.

2. Model Description

The model used here to simulate the performance of lithium-ion battery cells is based on a finite-element approach that defines the macroscopic geometry of the cell. The cell is divided into a collection of rectangular elements as shown in Figure 1, with five elements to cover the cell height, two elements to cover the cell width, and three elements (anode, separator, and cathode) to cover a layer in the electrode stack within the prismatic cell case. All layers in the prismatic stack, which can consist of several hundred parallel-connected anode/separator/cathode layers, are assumed to behave identically in this model. The model allowed the dimensions of the cell, the thickness of the internal electrodes and separator, and the overall plate surface areas to be adjusted to match that of the actual cell design being modeled.

The macroscopic assembly of elements in Figure 1 is used to model the movement of electrons, ions, electrolyte, and heat throughout the cell components in response to the electrochemical processes of charge and discharge. The movement of these species through the cell model is governed by a circuit that describes each flow rate in terms of the driving potentials and the resistances to flow from element to element. For example, electronic and ionic flow through the cell is governed by the circuit shown in Figure 2, where the electrochemical processes provide the driving potentials, and the circuit is composed of the charge transfer, and electronic and ionic resistances in the cell components.

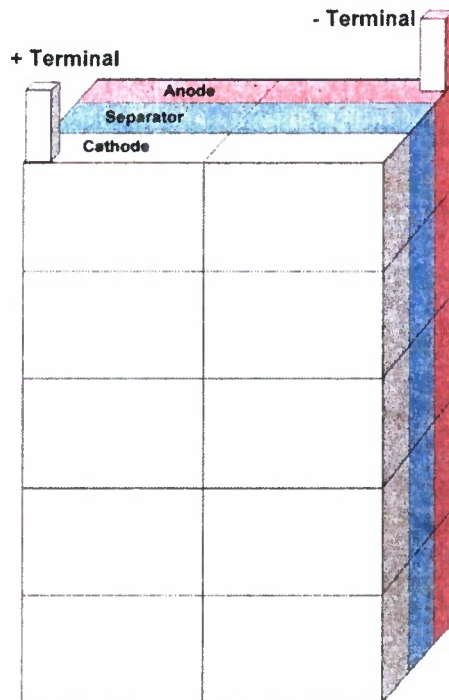


Figure 1. Finite elements used to model the macroscopic spatial dimensions of a prismatic lithium-ion battery cell.

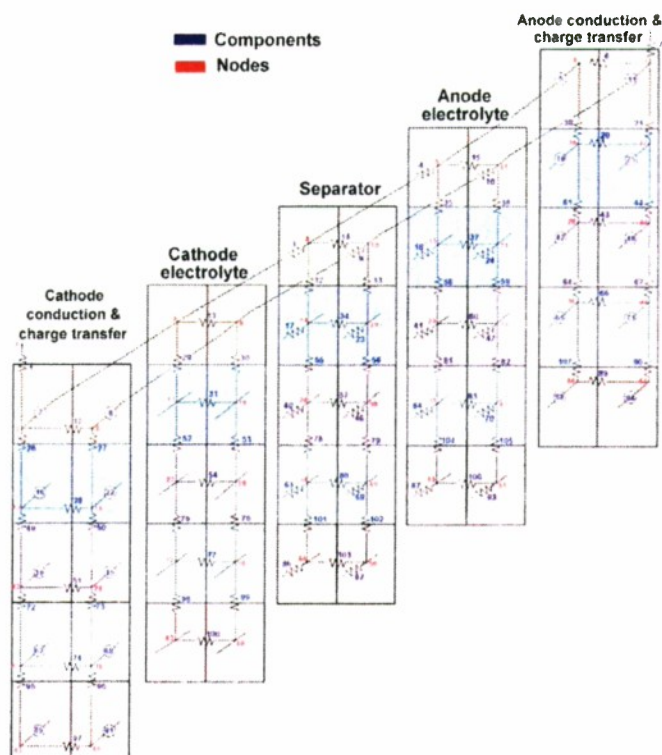


Figure 2. Electrical circuit showing electronic and ionic conduction, as well as charge transfer through the cell model. Out-of-plane connections between layers are only shown at the top of the cell.

Similar circuits are used to describe the diffusive flow of chemical species throughout the cell in response to concentration gradients and the movement of heat in response to thermal gradients across the cell.

The electronic resistances in the anode and cathode layers in Figure 2 were calculated from the dimensions and thickness of the anode and cathode foils used in making the cell. Anode foils were assumed to be made of copper, and cathode foils were assumed to be aluminum. The active material thickness coated onto the foils was based on the manufactured thickness for the electrodes along with the porosity of the coatings. Overall surface area was dictated by the number of electrode layers put into the cell by the manufacturer. The percentage of binders and conductive diluents in the active material coatings could be specified as needed for any particular cell design. The separator was assumed to be a microporous polymer layer that was specified by its thickness and its porosity.

Ionic conduction in the electrolyte was allowed to occur through the liquid electrolyte that was assumed to fill all the pore volume in the anode, separator, and cathode. A Li-ion transference number was used to specify the proportion of the current carried by the lithium ions in the electrolyte, as opposed to that carried by the electrolytic counter-ions. The variation of electrolyte conductivity was dynamically computed as the temperature and concentration changed within each element during cell charge and discharge. The conductivity characteristics of any electrolyte composition used by the manufacturer could be provided as model inputs.

Relatively complex models were used for the charge transfer components in the anode and cathode layers of Figure 2, which are represented in each of these layers by numbered circles. Much of the change that is likely to occur within the cell as it ages is in the details of the transfer of charge from electronic carriers in the foils to ionic carriers in the electrolyte. In the cathode, the charge transfer model involves five discrete processes, as illustrated in Figure 3.

1. Electronic conduction from the metal foil into the conductive diluent in the cathode, and then through this diluent material to the surface of an active material particle.
2. Electronic conduction from the conductive additive into the active material lattice and to a site in the lattice where charge transfer is to take place.
3. The electrochemical charge transfer process.
4. Diffusion of lithium ions through the metal-oxide active material lattice to the interface between the oxide particle and the electrolyte.
5. Movement of a lithium ion from the surface of the metal-oxide lattice into the liquid electrolyte.

The first of these cathode charge transfer processes is modeled as a series resistance having a value dictated by the amount and conductivity of the additive assumed to be in the cathode, and the proportion of the active material surface area in contact with the additive.




Figure 3. Illustration of the charge transfer processes modeled in the cathode active material.

The second cathode process is modeled as a series transmission line resistance in which electrons are “transmitted” into the interior of the metal-oxide particles according to the electronic resistance of the oxide lattice. For metal oxides having very low electronic conductivity, as is the case for the nickel and cobalt oxides commonly used in Li-ion cells, there is typically very little electronic conduction into the interior of the oxide particles. In this situation, the charge transfer sites are typically close to where the surface of the oxide particle contacts an adjoining electronic conductor.

The third cathode process is the electrochemical charge transfer step, which involves the reaction of the electron with a mobile lithium ion at a metal-oxide charge transfer site within the active material lattice. This process is modeled using the Butler-Volmer equation, and involves the anodic and cathodic exchange current densities, the reversible potential for the charge transfer reaction in the cathode as a function of lithium content in the lattice, the temperature, and the number of available charge transfer sites within that particular finite element.

The fourth cathode process involves the diffusion of lithium ions in the metal-oxide lattice from the site where they were electrochemically formed by charge transfer to a site at the surface of the metal-oxide particle where it contacts the liquid electrolyte. This process is assumed to be governed by standard diffusion kinetics over the average distance from the charge transfer site to the surface of the particles in that finite element. The rate of this process depends not only on the magnitude of the Li-ion diffusion coefficient, but on the state of charge of the particles and the size of the active material particles.

The final process in the cathode charge transfer sequence models the dynamics of transferring a lithium ion from the surface of a metal-oxide particle to a lithium ion dissolved in the liquid electrolyte phase. This transfer requires that the surface layer of the metal oxide become electrostatically polarized before facile ionic transfer can occur across the boundary. This process is illustrated by the diagram in Figure 4. All metal-oxide electrodes require this type of surface polarization to develop




Figure 4. Illustration of the polarization layer required at the cathode active material to enable ionic transfer to the electrolyte in Li-ion cells.

before ions in the solid phase can readily exchange with ions dissolved in the electrolyte. Since the polarization that is required during discharge has a sign opposite that required during charge, this process causes a capacitive voltage hysteresis between the charge and discharge behavior of the cathode. This process is modeled as a surface polarization potential that is formed by capacitive charging of a surface layer. The model parameters that control this process are the magnitude of the surface polarization potential and the surface capacitance associated with generating the surface polarization. When no current is flowing through the cathode, the surface polarization layer is assumed to gradually depolarize in response to thermal movement of ions through the layer.

The charge transfer model used for the anode is illustrated in Figure 5. Here lithium ions diffuse through an SEI layer at the surface of the carbon particles, and electrons undergo conduction to a charge transfer site where the lithium ion is reduced to an intercalated lithium atom.

The key processes modeled in this process are: the movement of lithium ions through the SEI layer, which is the SEI resistance in the model; the conduction of electrons to a charge transfer site, which is the anode conductivity in the model; and the charge transfer process, which is governed by an anode Butler-Volmer expression in the model. In this model, the lithium atom movement within the carbon lattice of individual carbon particles is assumed to be sufficiently facile that it does not limit the rates of the reactions. The average lithium content of the particles is used in the model, thus neglecting any state-of-charge gradients within individual carbon particles.

The model described above for a prismatic lithium-ion battery cell contains a large number of parameters that describe the processes in the anode, electrolyte, and cathode. Many of these parameters describe physical characteristics of the cell and electrode structure, or the fundamental electrochemical characteristics of the materials, and are not expected to change as a cell ages. These types of parameters include those related to the thicknesses and sizes of the components, the reversible




Figure 5. Illustration of the processes involved in the anode energy storage materials.

potentials and other electrochemical parameters of the anode and cathode materials, and the electrolyte conductivity characteristics. A number of other parameters could clearly change as cells age and undergo performance degradation. We have analyzed the parameters in this model and have identified nine parameters that are very likely to undergo significant change as the result of the degradation modes commonly expected in Li-ion cells with age. These parameters are:

1. Total cathode capacity for Li ions (when charged and held at 4.1 V vs. Li).
2. Initial cathode state of charge
3. Lithium ion diffusion rate in the cathode active material
4. Total anode capacity to reversibly intercalate lithium
5. Amount of charge remaining in anode or cathode when the cell is fully discharged (anode/cathode charge balance)
6. Charge transfer resistance in cathode
7. Cell series resistances (anode SEI and electrolyte resistance)
8. Cathode surface polarization rate
9. Cathode surface polarization potential

In using this model to describe the performance of a lithium-ion cell through its lifetime, these nine parameters were allowed to change to obtain the values that best matched the observed cell performance. All other parameters were fixed at the values appropriate for the cell and its components as it was built at the beginning of its life.

3. Genetic Adaptation to Test Data

The approach developed here used the cell model described above to determine the most likely processes responsible for degradation during the life testing of a lithium-ion cell. The general methodology involved fitting the cell performance during the life test to the model, with the model parameters that gave the best fit to the data indicating the nature of the chemical or physical degradation in the cell as it went through life.

Initial attempts to perform this analysis used standard regression techniques to fit the nine parameters described above to the test data, thus obtaining the best fit of the model to the data. Unfortunately, standard regression methods based on minimizing the deviation between the model prediction and the observed cell voltage were found to be incapable of providing a unique and optimized fit to the test data. The problems that prevented a successful regression analysis were twofold. First, strong interactions existed between the nine parameters described above in terms of how they affected the predicted cell voltage. This resulted in extremely slow convergence of the model to a particular best fit, and made the result quite sensitive to the initial values chosen for the adjustable parameters. Second, the model solutions contained a high density of local minima. The regression solution would invariably get trapped in one of these local minima and provide an apparent solution that was not the true solution. This could be recognized by finding a different “optimized” solution each time the regression analysis was run.

Because of these regression difficulties, an alternate method was sought to obtain the model solution that best described a cell current or voltage behavior at a given point in the lifetime of a cell. The method that was adopted utilized a genetic algorithm to adapt an entire population of models to the cell data. This genetic technique was found to give reproducible model solutions from independent runs with the same data, and these solutions typically displayed a 30–50% improved fit to the data than was obtained from any given standard regression analysis.

The basic concept behind the use of genetic adaptation for modeling physical systems is that a diverse population of solutions is continually allowed to adapt to the data. The values of the parameters that describe the physical system are the “genetic code” that makes each member of the population unique. For the model described above, the nine key parameters comprise the genome, and each population member has a different value for each of these nine genetic codons. The “fitness” of any individual model is proportional to how well it fits the measured cell data. Maintaining population diversity over a wide range of parameters keeps the population from falling into a “less-than-optimum” genetic dead-end, and allows the population of models to continually explore the full range of parameters available to the model architecture. Solutions that are less fit die off quickly, while solutions that are more fit live a longer life and propagate into the population.

The size of populations used in genetic analysis is important, but is generally not critical as long as the population size is sufficient to be statistically meaningful. Typically population size should be more than 30 individuals, with little benefit being offered by having more than 100 individuals.

Populations will typically develop by producing sequential generations of increasingly “fit” individuals. This occurs by genetic mixing between individual members to produce offspring. The longevity and thus the number of offspring produced by each population member are in direct proportion to that member’s fitness. Extremely fit individuals could have multiple offspring, while extremely unfit individuals may not survive long enough to produce any offspring.

Diversity, an important element of any population, is typically maintained by several mechanisms. The simplest of these is through genetic mutation, which consists of sudden random shifts in the genetic code of an individual. Although these shifts are frequently not beneficial in terms of improved fitness, they help maintain a diverse population. Extremely unfit mutations will of course die off rapidly without propagating into the population.

The final mechanism for genetic adaptation actually occurs over the lifetime of each population member in response to the environment that the individual experiences. This effect is termed “epigenetic drift,” and results whenever a mechanism is included for tweaking the genome slightly in response to either harsh or benign environments. In living organisms, this mechanism is epitomized by the methylation of cytosine bases in DNA, which can either activate or silence particular genes and their functions. In physical models such as battery cells, this mechanism could simply involve a slight “tweak” to the genome of an individual where an obvious fitness advantage could be gained. Whether or not these shifts in the genetic code are passed on to offspring depends on when the genetic algorithm causes them to occur during the lifetime of an individual.

3.1 Genetic Algorithm for Modeling Li-Ion Cells

A key first step in the genetic analysis of any physical system is to define the initial population of solutions that will serve as the starting point for the adaptations that occur during future generations of the population. In our lithium-ion cell model, the initial population consisted of 40 individual models. This initial population was chosen to be the fittest 40 individual models from a randomly generated population of 400 models. These 400 models were obtained from random permutations of the nine adjustable model parameters over the entire range of physically reasonable values that each of them could adopt. For example, cathode capacity early in the life of a 70-Ah lithium-ion cell might be constrained to fall between 55 and 85 Ah, and initial models using random values in this range would be allowed.

Once the initial population of 40 individual models is obtained, each individual is allowed to live out its lifetime, with the quality of that individual’s environment affecting its genome, and the individual living a lifetime and having offspring in proportion to how well it fits the data. Each generation is marked by the replacement of a member of the population with an offspring model. Since only the 40 fittest models are retained following the lifetime of all the individual models, the least fit offspring may not be genetically represented from one generation to the next.

The genetic algorithm used here to produce offspring models utilizes a linear mixing of the genetic code (model parameters) for two models. These two models are selected using roulette probability criteria; i.e., the probability of being selected is in direct proportion to the model fitness. The fitness of the offspring is calculated, and if the offspring model is less fit than the parent models, it is dis-

carded and not included in the next generation. If the offspring model is more fit than the parents, it is included in the next generation of the population. The linear mixing of the genetic codes to produce offspring includes a diversity index that allows offspring genetic mixtures up to a prescribed amount outside the range of that parameter for either parent. This was found to be a critical diversity factor in allowing a population to shift parameters outside the ranges defined by the initial population if such a shift was found to achieve improved fitness. A diversity index of 30–50% gave the adaptation algorithm optimal flexibility, while significantly higher diversity indices could prevent the population from converging to the optimum solution.

The algorithm used here for mutations consisted of a simple shift in a randomly selected model parameter for an offspring model. These random mutations occurred at a specified rate for the offspring models, and were constrained to be within the bounds that specified the initial model parameters. If random mutations occur at an excessive rate, typically more than several percent of the time, they can prevent the population from converging to an optimum solution. In the genetic algorithm used here, a 0.5% mutation rate was used. This was found to maintain genetic diversity while not preventing the population from converging to the best fit to the data.

Epigenetic drift in the genome of each individual model was also included during the lifetime of that model. This process allowed the parameters describing the genome of each model to shift gradually during the model's lifetime in response to the model's environment. The environment E for each genomic parameter δ was assumed to be $E = dF/d\delta$ with all other parameters held constant, where F is the quality of the fit to the cell data. Drift in the model parameters was only allowed if it led to an improvement in fitness F , and the drift in each parameter was allowed to proceed until the fitness was maximized (deviation from data was minimized) for that parameter. This epigenetic drift was quite significant for models that were far from a local or global minimum in F , but was negligible when the model population had converged to the global minimum that defined the optimum fit to the data.

The factors described above comprised the rules that governed the changes in model parameters as one model population was replaced by another population from generation to generation. The actual sequence by which these rules were applied was as follows:

1. A starting population was randomly selected.
2. The population members were allowed to undergo epigenetic drift according to their local environments.
3. Offspring models were computed from population pairs, including periodic mutations.
4. The population was allowed to continue to produce offspring until 40 new models having better fitness than their parents had been created.
5. The fittest 40 models were used to seed the next generation.
6. Steps 2 through 5 continued until the improvement in average fitness for the fittest individual model over four complete generations improved by 0.000002 V or less.

Genetic algorithms often allow a select group of elite individuals (those being the most fit) to survive from one generation to the next. This prevents the model having the best fit from being discarded,

which is important if no child models having a better fit can be found. The danger in allowing elite individuals is that they can coerce the entire population into a local minimum that is not optimum since they are not subject to adaptive selection. In the algorithm used here, the fittest individual was only allowed to survive into the next generation if it could produce no improved offspring models after producing an average of two offspring with all other population members (this set an upper allowed limit of 80 total offspring). This was seen to occur very infrequently, even after convergence of the population to the global optimum. Typically, a given individual produced less than 10 offspring during each generation, with the actual number varying according to fitness and random selection.

3.2 Test Data Used for Modeling

The test data used to derive the model parameters that best fit the observed performance of a selected lithium-ion battery cell was taken from a life test that is presently underway for a prismatic Li-ion cell that had approximately 70 Ah of capacity when fully charged at the beginning of its life. This cell was stored at ambient temperature for 18 months, and was then placed into a life test that involved cycling the cell to 40% depth of discharge based on its rated capacity of 56.66 Ah. Each cycle was 96 min in duration, with 35 min for discharge and the remaining time for recharge. Recharge was done using a peak charge rate of 25 A, and once the cell voltage reached 4.0 V, the current was tapered so as to hold the cell at 4.000 V during the remaining recharge time. Figure 6 shows the voltage and current from a typical 40% depth-of-discharge cycle during this life test.

The base plate on which the cell was mounted was held at 20°C, and the cell temperature was allowed to warm and cool according to the balance between its heat generation rate and the cooling from the base plate. No convective cooling was allowed from the cell case to the atmosphere surrounding the




Figure 6. Typical voltage and current behavior during a 40% depth-of-discharge cycle.

cell since the cell case was wrapped in a ceramic wool blanket. Each broad face of the prismatic cell was clamped to an aluminum plate that could conduct heat to the base plate. The temperature was measured at the top of the cell as well as on the base plate near the base of the cell. Figure 7 shows a typical thermal profile for this cell during the 40% depth-of-discharge cycling.

About every 2000 cycles during this life test, the cell was discharged at the cycling discharge rate (38.9 A) to 3.0 V to measure its usable capacity. This technique was used to evaluate how much capacity margin remained after the end of the 40% depth-of-discharge cycle at that point in the life-time of the cell. The cell was then recharged at C/5 to the 4.00-V limit and put back on cycling at 40% depth of discharge. The C/5 recharge rate was used because that was approximately the rate to which the current tapered down at the end of charge during the repetitive 40% depth-of-discharge cycling. The voltage during charge and discharge for these periodic capacity checks is shown in Figure 8.

While it is clear in Figure 8 that some loss of capacity is occurring as the cell ages, the changes in cell resistance are not quite as clearly defined. During discharge, the lowered voltage makes it appear that the cell resistance has increased significantly, but during recharge, there is no obvious indication of the elevated charge voltage expected when the cell resistance increases. The capacity cycles shown in Figure 8 comprise the cell voltage database (with cell temperature and current) that was used to determine the model parameters that best describe the changes to date in this cell over its life.

The voltage spike in Figure 8 at 40% capacity during discharge resulted from a 20-s change in current that was used to provide an indication of changes in the resistance of the cell at its cycling depth-of-discharge point. The modeling done in this report did not directly use this current change to determine cell resistance, but rather obtained cell resistance changes by fitting the complete charge/discharge behavior of the cell (including the 20-s current change) to our model.




Figure 7. Typical cell and base plate thermal profile during the 40% depth-of-discharge cycling.




Figure 8. Capacity cycles used to model Li-ion cell changes during life testing.

3.3 Adaptive Modeling Approach

The lithium-ion cell model described earlier was allowed to adapt to a given set of test data using the genetic algorithm described earlier in this section. Starting with the initial dataset in Figure 8 (at 500 cycles), the nine model parameters described earlier and shown in Table 1 were genetically adjusted to provide the optimum model fit to the data. These nine parameters were initially assumed to fall randomly between the upper and lower limits given in Table 1. The initial limits for the parameters were only adjusted up or down if the preceding optimized parameter was found to drift toward one of the limits. For example, the initial cathode capacity was initially assumed to fall between 50 and 70 Ah, and the optimized value at beginning of life was nearly centered in this range at 59.89 Ah. However, this parameter had drifted down to 55.91 Ah after 7563 cycles; thus, an initial range of 46 to 66 Ah was used at 7563 cycles. Similarly, the upper initial limit for the Li^+ diffusion rate was reduced as

Table 1. Initial Adaptation Ranges and Optimized Values for Parameters Adjusted in the Cell Model to Describe Changes in Cell Performance During Life Test

Parameter	Initial parameter range		Optimized values at cycle				
	Minimum	Maximum	500	1500	3496	5505	7563
Initial cathode capacity (Ah)	50	70	59.89	59.503	58.127	57.473	55.91
Total cathode capacity (Ah)	60	80	70.146	70.909	69.474	69.964	68.47
Li^+ diffusion rate ($D \cdot \text{area}/\text{length}$)	0.03	0.12	0.1092	0.07523	0.06103	0.05483	0.05233
Surface polarization time const (per coul cm^2)	0.1	1	0.471	0.3575	0.2403	0.2084	0.2968
Cathode surface polarization potential (V)	0.005	0.03	0.0204	0.0215	0.0269	0.0302	0.0354
SEI/electrolyte conductivity (per ohm cm)	1	50	20.61	54.64	33.81	61.43	7.777
Total anode capacity for Li ions (Ah)	62	82	73.83	72.864	73.25	72.772	73.45
Anode precharge (Ah)	-10	5	-3.775	-3.424	-5.031	-4.577	-6.054
Cathode charge transfer resistance (ohm cm^2)	0.5	4	1.079	1.303	1.379	1.569	1.73

the diffusion rate was found to drift downward over life. Table 1 also lists the optimized values for the nine parameters that best fit each of the datasets illustrated in Figure 8 obtained during the first 7563 cycles of this cell life test.

Figure 9 shows the initial and optimized populations of 40 models for the first and third parameters listed in Table 1 at 5505 cycles. The optimized population of 40 models has a sufficiently small variability in these two parameters that it appears as a single point on the scale of Figure 9. The adaptive optimization that took place to get from the initial population to the optimized population in Figure 9 involved the birth and death of 3360 generations of models.

The process of genetic adaptation to a particular lithium-ion cell dataset typically involves adaptive changes over many thousands of generations of models. A typical example is shown in Figure 10. The genetic adaptation method typically results in a relatively rapid decrease in average deviation from the data in the first several hundred generations as the population converges to the best model within the range of initial parameters. The example shown in Figure 10 also includes a slow convergence over thousands of generations to the ultimate global best fit to the data. This slow convergence resulted because, for this example, one parameter (cathode capacity) was intentionally given an upper




Figure 9. Distribution of initially selected and optimized populations for two of the nine parameters adjusted to model cell performance changes over life.




Figure 10. Changes in average deviation between the model and the data during the 3360 generations involved in the genetic optimization illustrated in Figure 9.

initial capacity limit that was nearly 10 Ah lower than the correct cathode capacity. The purpose of this was to determine whether (or how efficiently) the adaptive algorithm could find a global best solution that was well outside the initially specified parameter range. The capability to search for and find an optimized solution outside the initially specified range of parameters is enabled by epigenetic drift and by mutations. The result, as shown in Figure 10, was that the genetic method, using a diversity index of 50%, was capable of reproducibly converging to a solution that could be well away from the anticipated parameters. This capability is critical for a robust modeling method since battery cell data can frequently be influenced by unexpected changes in cell parameters.

For the datasets analyzed in this study, typical average deviations for the optimized models were between 0.003 and 0.006 V. Figure 11 shows a typical optimized fit to the test dataset (at 5505 cycles), and illustrates that the optimized model was capable of fitting the measured test data quite well.




Figure 11. Measured and optimized model cell voltages after 5505 cycles.

4. Analysis of Model Results

The optimized model parameters listed in Table 1 can be used to determine how important each of the processes included in this Li-ion cell model is in contributing to the degradation of cell performance as this cell ages. We have done this by trending the changes in the nine parameters in Table 1 as a function of cell life. The following sections provide the results of these analyses, as well as a final section that uses the trends in conjunction with the performance model to predict how rapidly cell performance is expected to degrade in the future. By periodically updating this prediction of cell lifetime as the life test progresses, we can evaluate how effectively the early life data capture the changes that will eventually cause the cell to fail. If future life predictions remain relatively consistent with these results, our confidence in the fidelity of the life prediction will increase, while significant shifts in future life predictions would indicate that the process responsible for ultimate cell failure may not be adequately captured in the modeling performed here.

4.1 Initial Cell State of Charge

The initial state of charge of a cell that is undergoing rapid cycling can be an extremely important factor in contributing to cell failure. This is because the rapid cycling never allows the cell to completely stabilize at the full charge possible at its end-of-charge voltage. This is recognized by the charge current not having an opportunity to fully taper down to levels close to zero at the end of charge. The main factor contributing to decreases in state of charge as the cell starts each discharge cycle is increasing cell impedance, which tends to make the cell current begin tapering earlier, and results in a slower current taper during the remaining charge time. The cell can ultimately fail if the initial state of charge is insufficient to provide the required discharge capacity, and failure onset can occur very rapidly if the cell resistance increases to the point where cyclic energy balance cannot be maintained.

Figure 12 shows the trend in initial cathode state of charge from the model computations. The results show a linear decrease in the cathode state of charge to date.




Figure 12. Model trend in initial cell state of charge during cycling.

The trend shown in Figure 12 only reflects the active lithium capacity available in the cathode at the beginning of discharge, and thus its changes could reflect change in active or inactive lithium content. At this point in the life of this cell, there are no data suggesting that changes in initial cathode state of charge will significantly depart from linear behavior in the future, thus a linear trend was used for this parameter over the cell lifetime. This trend will be adjusted in the future if additional data show a departure from linearity.

4.2 Cathode Capacity

The cathode capacity is a parameter that indicates the total capacity of the cathode to store lithium to a reversible voltage of 4.1 vs. Li. Thus, changes in this parameter would indicate either that the cathode structure is undergoing degradation to make the active material less capable of storing lithium, or some of the cathode material has changed to a form that does not reversibly store lithium. Figure 13 shows the trend observed to date in this parameter.

Again, the data suggest that so far the total cathode capacity has only decreased about 2 Ah, and that this decrease is essentially linear with cycling. While the results in Figure 13 make it appear that there is a large scatter in the points for this parameter, the scatter is largely because there has only been a small change in the total cathode capacity to date. As further cycle life data are obtained, it may be possible to detect departures from linearity, but at present, this parameter will also be assumed to decrease linearly through the cycle life of this cell.

4.3 Lithium-Ion Diffusion Rate

The rate at which lithium ions can diffuse through the cathode active material is recognized as key to the performance of Li-ion cells. Because the cathode active material has only limited electronic conductivity, access to charge stored within the cathode lattice is largely enabled by lithium-ion diffusion through the lattice. Here, we define the lithium-ion diffusion rate as the diffusion coefficient times




Figure 13. Model trend in total cathode capacity for lithium during cycling.

the ratio A/L , where A is the average cathode active material area over which diffusion can occur, and L is the average distance for diffusion. To trend the changes in diffusion rate, the diffusion resistance (the reciprocal of diffusion rate) has been plotted as a function of the square root of the number of cycles. The square root dependence was selected because the data could not be accurately described using a linear dependence on cycle life. This plot is shown in Figure 14, where the points are the model results based on the measured data. There appear to be significant changes in the diffusion rate during the cycling to date, although it does appear to be stabilizing somewhat. The solid line is an empirical function that was found to fit the data points presently available quite well.

Interestingly, the diffusion resistance was not found to change linearly with the square root of the cycles, as would be expected from a simple layer growth model. It has become common practice in the business of modeling lithium-ion cell degradation to include a degradation component proportional to the square root of the cycles since this has often been found to empirically describe the observed overall cell changes. The results in Figure 14 suggest that there may not be any particular degradation processes following a linear square root of time dependence, but that the frequently observed empirical dependence simply results from the averaged degradation rates of a number of cell processes.

The empirical solid curve in Figure 14 follows the form $y = A\ln(x) + B$. Clearly, the form of this curve is arbitrary since it is not based on fundamental first principles, and it probably will require adjustment as more data become available. The key issue is probably whether the curve will continue upward with some significant slope, or whether it will level off and stop changing as the cell ages. However, at present, the indicated curve accurately describes the observed trends, and will be used here for projecting how the lithium-ion diffusion rate is expected to change in the future.




Figure 14. Model trend in the lithium-ion diffusion resistance during cycling.

4.4 Cathode Surface Polarization Time Constant

The surfaces where the cathode active materials contact the electrolyte develop a polarized surface layer as part of the process of transferring lithium ions from solution into the solid active material. This surface polarization process has a time constant that is proportional to the number of coulombs per cm^2 that are required to fully polarize the surface. This time constant does not tend to have a large effect on cell performance unless it becomes relatively large (more than 0.5 to 1 coulomb per cm^2). The polarization time constant as a function of cycles is indicated in Figure 15, which shows that it has decreased by roughly a factor of 2 during the cycling to date. At present, this parameter appears to have leveled off, although added data will indicate whether the increase seen for the last point continues. At present, the changes in this parameter are modeled by the solid curve as a monotonically decreasing function that levels off as the cycling continues.

4.5 Cathode Surface Polarization Voltage

The cathode surface polarization process described above results in an electrostatic voltage separation between the solid cathode lattice and the liquid electrolyte. This polarization decreases the discharge voltage and increases the charge voltage by the voltage needed to pass lithium ions from the solid to the electrolyte. At beginning of life this polarization, which shows up as a voltage hysteresis between charge and discharge, is 10–20 mV in magnitude. As a cell ages, it is not unusual for the surface polarization of the cathode to increase significantly. This change can be a significant factor contributing to cell degradation and ultimate failure.

The changes in cathode surface polarization voltage obtained from fitting the measured life test data to the cell model are shown in Figure 16. This parameter shows essentially a linear increase with time. Based on these data, we assume at present that the surface polarization voltage will continue to increase linearly through the life of this cell. This is clearly a significant factor contributing to cell degradation if it continues to increase at the 0.02 V per 10,000 cycles seen in Figure 16.




Figure 15. Model trend in the cathode surface polarization time constant during cycling.




Figure 16. Model trend in the cathode surface polarization voltage during cycling.

4.6 SEI/Electrolyte Resistance

The components in the cell that contribute fixed series dc resistances, such as the leads, electrolyte, and anode SEI (solid-electrolyte-interphase) were all lumped together into a single series-resistance parameter for modeling. The reason for this is that all these series resistances produce identical current/voltage behavior (Ohm's Law behavior), and thus cannot be readily resolved from each other by simply fitting the cell voltage and current for a given cycle. These parameters could potentially be separately determined by fitting cell behavior at differing temperatures. However, it is also expected that the electrolyte resistance and the resistance of the leads should not change substantially as the cell ages, while the SEI resistance is very likely to undergo changes as the cell is cycled. For this reason, changes in this series-resistance parameter are primarily attributed to changes in the anode SEI. Figure 17 shows the change in the SEI portion of the series resistance that was obtained for this cell during its cycling.

Upon first inspection of Figure 17, it appears that there is a quite large scatter in the values determined for the SEI resistance parameter. In actual fact, the resistance parameter produces a very small contribution to the overall cell polarization for any of the values in Figure 17. Thus, the large variation in Figure 17 appears as scatter and does not display any systematic variations with cell age. While such systematic variations may appear in the future for this cell, at present, it is assumed that the SEI resistance parameter is constant at a value given by the average of the points shown in Figure 17 for all projections of cell performance in the future. Because the impact of the SEI on cell performance is small, at present, this assumption has little effect on the cell voltage or current behavior.

4.7 Anode Capacity

The total capacity of the anode to reversibly store lithium is defined here as the anode capacity. This is not to be confused with the actual state of charge of the anode, which is the percentage of the total




Figure 17. Model trend in the SEI resistance (ohm cm^2) during cycling.

lithium capacity that has actually been charged with active lithium. The anode capacity obtained from modeling the charge/discharge curves of this cell is shown in Figure 18. This parameter appears to be somewhat scattered until one notes the quite expanded vertical scale.

The total anode capacity appears to change very little during the cycling of this cell, and the change that is seen is assumed to be linear with cycles. Since the slope of the linear trend is only -0.2 Ah per 10,000 cycles, this term does not have a significant effect on cell degradation based on the data to date.




Figure 18. Model trend in the total anode capacity for lithium during cycling.

4.8 Anode Precharge

The anode precharge is defined as the amount of active anode capacity that remains after the cathode is completely discharged, and thus, may have either a positive or a negative value. This parameter defines the capacity balance that exists between the anode and the cathode. A positive value represents the active capacity remaining in the anode after the cathode has been completely discharged (cathode limited cell). A negative value represents the active capacity remaining in the cathode after the anode has been completely discharged (anode limited cell). A cell can start out as cathode limited early in life, and then transition to anode limited at a later point in its lifetime as various chemical oxidation or reduction processes take place within the cell to shift the balance between the electrodes. Typically, net chemical oxidation of components within the cell will result in an increase in anode precharge, while net reduction will result in a decrease in anode precharge. Thus, this parameter can provide an indication of the types of parasitic chemical degradation processes taking place as a cell ages.

The changes in anode precharge for this lithium-ion cell are indicated in Figure 19 as a function of cycles. The anode precharge is negative, indicating an anode-limited cell, and appears to be decreasing linearly about 3 Ah every 10,000 cycles. This behavior suggests that a net chemical reduction of the internal cell components must be occurring to offset the oxidation associated with the decreasing anode precharge parameter. The linear trend shown in Figure 19 was used for making projections of cell performance based on the present data.

4.9 Cathode Charge Transfer Resistance

The charge transfer resistance within the cathode was assumed to be a resistive parameter that followed the current/voltage behavior specified by the Butler-Volmer Equation. At low currents, the current/




Figure 19. Model trend in the anode precharge during cycling.

voltage behavior is linear, while at high currents it is exponential. For the currents employed in this life test, the cell operated in the linear range of the Butler-Volmer Equation. The charge transfer resistance, plotted in Figure 20, displays a linear increase with cycling. The linear trend in Figure 20 was used to project how this parameter is expected to change in the future as the cell continues to cycle.

4.10 Life Projection Based on Trend Analysis

The projections of how the various parameters will change in the future as this cell continues to cycle can be used to predict the cycle life of this cell under its present conditions of operation. At any point in the future of the cell, the parameters that describe its performance in the cell model can be obtained by extrapolation of the observed trends shown in the preceding sections. This analysis has been done as a function of the number of cycles until the end-of-discharge voltage of the cell falls to 3.0 V, whereupon the cell is assumed to fail under these operating conditions. The results of this analysis are shown in Figure 21.

The red curve in Figure 21 shows the capacity margin remaining above 3.0 V, the blue curve shows the end of discharge voltage predicted by the model, and the blue squares indicate cycling data to date. The model predicts that the cells will fail after 42,050 cycles if cycling is continued under the present conditions. In actual practice, this test is designed such that the peak recharge voltage will be increased if the capacity margin falls below 11.33 Ah (20% of rated capacity). Thus, the model predicts that after about 35,000 cycles it will become necessary to start increasing the peak recharge voltage. Such increases in peak charge voltage (up to 4.1 V maximum) may allow an added 7,000 cycles of life to be obtained from the cells if degradation rates do not increase significantly as a result of the increase in peak charge voltage.




Figure 20. Model trend in the cathode charge transfer resistance during cycling.




Figure 21. Model projection of cycles until the end-of-discharge voltage falls to 3.0 V after 42,050 cycles. The squares represent data obtained to date.

Clearly the model does a good job of fitting the data obtained to date. The question of greater interest, however, is how well the model captures the rates of the degradation processes that will ultimately limit the life of this cell. This is best evaluated empirically by making updated model projections every 6 months. If each updated model projection does not change significantly as more life data are included in the trending, a high confidence in the life projection will eventually be attained. On the other hand, if each updated projection shows a systematic shift, that shift can be used to both estimate the uncertainty in the model projections and potentially correct the model for processes that may not be fully captured in the present model. In either situation, additional test data will allow the model projections to be refined significantly. Using this approach, it is expected that an accurate and validated model prediction of life should be possible well before the life test has been completed.

Another interesting point is the comparison of this modeling approach to other methods for predicting cell life that are commonly used in the battery industry. Numerous data-trending models have been published that note that the capacity remaining in a cell that undergoes rapid cycling appears to decay at a rate that is linearly proportional to the square root of the number of cycles (or time). The fundamental physics and chemistry that underlies such dependence is not clear, however, since it is very probable that a number of degradation processes are responsible for the cell degradation in most cases. The capacity data from this 50-Ah cell may be similarly plotted as a function of the square root of the number of cycles, providing the result shown in Figure 22. This extrapolation method effectively ignores degradation from all kinetic effects and resistance changes, and attributes cell failure to loss of capacity. As is typically reported from other life tests throughout the battery industry, the capacity data appear to be accurately represented by a linear trend when plotted this way. However, the cycle life predicted from this trend corresponds to about 175,000 cycles. This cycle life is vastly different from the 42,050 cycles predicted by the first principles modeling technique discussed here. It will require additional data to show which of these methods is the more accurate. However, until additional long-term data are gathered, it is probably safer to stay on the side of the more conservative prediction.




Figure 22. Plot of the remaining cell capacity as a function of the square root of the number of cycles completed at 40% depth of discharge. The blue line indicates when failure is expected after the remaining capacity drops below 40% of the rated capacity, and corresponds to 175,000 cycles.

An alternative prediction method that has been used frequently over the years consists of a simple linear extrapolation of the test data out to where the capacity drops below what is required to support cycling at 40% depth of discharge. Again, this extrapolation method ignores all degradation from cell resistance changes. The reason this method is not favored by Li-ion optimists in the industry is that it typically predicts a cycle life that is substantially lower than that predicted by models containing square root of time dependences. However, it is instructive to determine how well the existing data from this cell are fit by a simple linear capacity degradation model, and how such a prediction compares with that obtained here by applying a first principles model to the test data. Figure 23 shows a linear fit to the cell capacity measurements, and indicates that the capacity data points can indeed be




Figure 23. Capacity trend when plotted as a linear function of cycles.

represented just as well by a linear trend as by the square root of cycles plot in Figure 22. However, a key difference is that the linear plot of Figure 23 extrapolates to a predicted cycle life of 44,070 cycles, which is relatively close to the life of 42,050 cycles predicted by our first principles model. Again, additional data will help resolve the question of model fidelity and how a state-of-the-art first principles model stacks up against simplistic extrapolations commonly used in the battery industry.

5. Conclusions

A state-of-the-art first principles lithium-ion cell model has been used to describe the changes in the performance of a 50-Ah cell during the first 7563 cycles of its life. This model was fit to the life test database using a genetic adaptation algorithm that provided degradation trends for nine different parameters that described the processes most likely to undergo change during the life of the cell.

The major process contributing to cell degradation seems to be a decrease in the lithium capacity that can be reversibly stored in the cell, most of which correlates with a decrease in the anode precharge. Thus, the capacity loss is primarily controlled by a shift in the balance between anode and cathode capacity that makes the anode prematurely undergo depletion relative to the cathode. The lithium that is not getting charged into the anode is moving into the cathode as the charge balance between the electrodes shifts. From these data, it is not clear what the origin of the irreversible reduction process is in the cell that is causing the gradual shift in lithium charge balance.

It is also likely that the increasing lithium diffusion resistance and increasing surface layer polarization on the cathode are contributing to the diminished chargeability of the cell, along with the observed decrease in cathode capacity. Thus, the cathode appears to be changing more than the anode during the cycling. However, the majority of the capacity degradation arises from the net loss of lithium that can reversibly cycle between the anode and cathode.

The rates for the nine processes contributing to the cell degradation were used to predict the amount of degradation from each process through the entire life of the cell. The performance of the cell was predicted as a function of the cell life using the first principles model, thus allowing the cycle life of the cell to be predicted. This prediction yielded a predicted life of 42,050 cycles. This value was slightly lower than the 44,070 cycles predicted from linearly extrapolating the cell capacity loss, and was much lower than the 175,000 cycles predicted by extrapolation the cell capacity loss on a square root of time scale. At present, the cycle life predicted by our model appears to be a reasonable estimate that must await additional data to be verified.

PHYSICAL SCIENCES LABORATORIES

The Aerospace Corporation functions as an “architect-engineer” for national security programs, specializing in advanced military space systems. The Corporation's Physical Sciences Laboratories support the effective and timely development and operation of national security systems through scientific research and the application of advanced technology. Vital to the success of the Corporation is the technical staff's wide-ranging expertise and its ability to stay abreast of new technological developments and program support issues associated with rapidly evolving space systems. Contributing capabilities are provided by these individual organizations:

Electronics and Photonics Laboratory: Microelectronics, VLSI reliability, failure analysis, solid-state device physics, compound semiconductors, radiation effects, infrared and CCD detector devices, data storage and display technologies; lasers and electro-optics, solid-state laser design, micro-optics, optical communications, and fiber-optic sensors; atomic frequency standards, applied laser spectroscopy, laser chemistry, atmospheric propagation and beam control, LIDAR/LADAR remote sensing; solar cell and array testing and evaluation, battery electrochemistry, battery testing and evaluation.

Space Materials Laboratory: Evaluation and characterizations of new materials and processing techniques: metals, alloys, ceramics, polymers, thin films, and composites; development of advanced deposition processes; nondestructive evaluation, component failure analysis and reliability; structural mechanics, fracture mechanics, and stress corrosion; analysis and evaluation of materials at cryogenic and elevated temperatures; launch vehicle fluid mechanics, heat transfer and flight dynamics; aerothermodynamics; chemical and electric propulsion; environmental chemistry; combustion processes; space environment effects on materials, hardening and vulnerability assessment; contamination, thermal and structural control; lubrication and surface phenomena. Microelectromechanical systems (MEMS) for space applications; laser micromachining; laser-surface physical and chemical interactions; micropropulsion; micro- and nanosatellite mission analysis; intelligent microinstruments for monitoring space and launch system environments.

Space Science Applications Laboratory: Magnetospheric, auroral and cosmic-ray physics, wave-particle interactions, magnetospheric plasma waves; atmospheric and ionospheric physics, density and composition of the upper atmosphere, remote sensing using atmospheric radiation; solar physics, infrared astronomy, infrared signature analysis; infrared surveillance, imaging and remote sensing; multispectral and hyperspectral sensor development; data analysis and algorithm development; applications of multispectral and hyperspectral imagery to defense, civil space, commercial, and environmental missions; effects of solar activity, magnetic storms and nuclear explosions on the Earth's atmosphere, ionosphere and magnetosphere; effects of electromagnetic and particulate radiations on space systems; space instrumentation, design, fabrication and test; environmental chemistry, trace detection; atmospheric chemical reactions, atmospheric optics, light scattering, state-specific chemical reactions, and radiative signatures of missile plumes.